SNARE: Social Network Analysis and Reasoning Environment

Doug Riecken*, Anita Raja+*, Rebecca J. Passonneau*, and David L. Waltz*
Columbia University Center for Computational Learning Systems*
University of North Carolina Charlotte+

"to learn something well, you must learn it many different ways"
- Marvin Minsky

Abstract

The importance of diversity in reasoning and learning to successfully address complex problems is examined. We discuss an approach by which a multiagent framework with decentralized control mechanisms provides diverse perspectives and hypotheses addressing a class of complex problems. We introduce the SNARE multiagent system. SNARE performs tasks to gain situational awareness of situations of interest in a Social Media Space. It applies a decentralized control mechanism for each agent; this mechanism enables an agent to interact with other agents to reason and learn. This approach facilitates dynamic agent organizations that adapt the topologies of interactions between agents based on the problem context.

Research Statement

We are interested in complex problems spaces that require multi-strategy reasoning and learning techniques to effectively identify and respond to emerging situations of interest. Our approach investigates the utility of independent reasoning agents, each agent being a distinct expert; these agents collaborate in various groupings.

Specifically we are investigating formal representations for classes of problems that fit and are tractable to leverage Decentralized Partially Observable Markov Decision Processes (DEC-POMDPs) to enable multi-strategy reasoning and learning by collaborating sets of agents. Further, we are examining formal representations that will allow extending the complexity and scale of various problem spaces that our reasoning/learning techniques address.

We will discuss our work on SNARE: Social Network Analysis and Reasoning Environment. SNARE is a distributed computing system of collaborating human analysts and independent specialized intelligent software components (a.k.a. software agents). Our current investigation involves leveraging SNARE in the domain of social media to obtain anytime situational awareness of social media spaces; a social media space is an excellent candidate problem space due to the fluid nature of its complex domain features.

Motivating Discussion

Complex problems that require expressive problem solving strategies suggests: (1) the utility of gaining multiple perspectives of a given problem solving situation, (2) formulating and evolving one or more hypotheses of the situation, and (3) providing one or more explanations of each hypothesis. The need for a diversity of reasoning/learning techniques and knowledge representations when addressing complex problems continues to be an important research topic [McCarthy, 2002].

We will consider a sequence of work that provides the grounding of our control and reasoning mechanisms in SNARE. First, prior work on the M system [Riecken, 1994] provides an example of collaborating agents that generate multiple perspectives, hypotheses, and explanations of a problem space. Each agent is a complex expert in a specific reasoning modality (e.g., temporal, causal, etc…). M demonstrated an expressive performance of its agents collaborating via a centralized blackboard control mechanism.

Next, work on the RESIN system [Yue, 2009] will be discussed to highlight extending the M work by applying a centralized Markov Decision Process (MDP) [Bertesekas, 2006] to implement RESIN's blackboard control mechanism. This work lays the groundwork to advance the use of MDPs as an approach towards a decentralized control mechanism that would enable each agent with its own ability to manage collaboration with other agents versus a centralized controlling architecture.

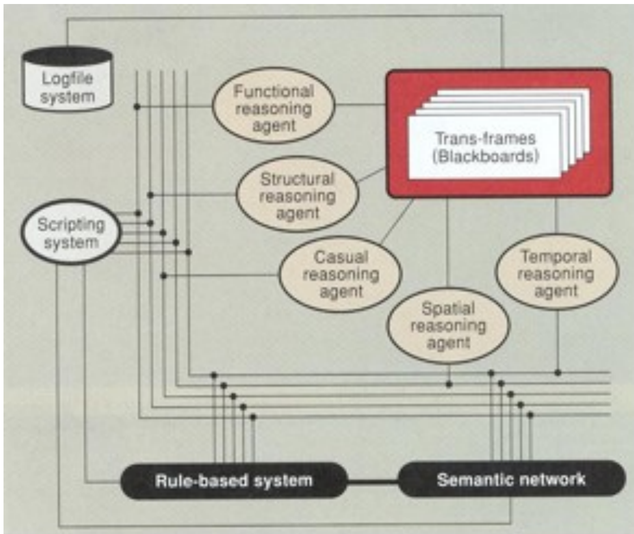

Figure 1: Architectural Diagram of the M architecture

Finally, we highlight our use of DEC-POMDPs and the SNARE system. We explain our formalized representation of DEC-POMDPs and how they fit our applied work in social media. The SNARE architecture leverages decentralized collaborating agents as an effective approach to: (1) reason from multiple perspectives, (2) provide a means by which each agent manages it reasoning and learning in collaboration with other agents, and (3) create a dynamic technique for agents to join and exit "agent communities".

M System

M was developed to support human interactions in a Computer Supported Collaborative Work (CSCW) environment referred to as a Virtual Meeting Room (VMR); M was developed at Bell Labs Research for an AT&T virtual meeting service. M served to provide personalized support for each person working in a VMR. It observed, identified, and classified the actions performed and artifacts created by persons working on a team in a VMR in order to hypothesize about situations in the VMR.

Since VMR situations are feature rich, M was designed with distinct modal expert agents. Each agent is a specialized modal reasoner; several examples include: temporal reasoner, spatial reasoner, structural reasoner, functional reasoner, and causal reasoner. As actions occur in a VMR, one or more reasoners would identify and hypothesize on the situations occurring. A temporal reasoner might identify and hypothesize that a set of actions and corresponding artifacts are related based on the temporal sequencing of the actions. Further, the spatial and functional reasoners may subsequently complement the temporal reasoner by elaborating on this situation of actions/artifacts as occurring in a specific location in the VMR and the artifacts all serve a similar function, thereby elevating the hypothesis classifying the actions/artifacts as being related. Each modal agent was engineered with its specific modal knowledge of VMR actions and artifacts. Also, over time, each agent learned the personalized behaviors of each user's typical performance and needs when using a VMR.

The agents collaborated as blackboard knowledge sources via a set of blackboards managed by a centralized blackboard controller. Blackboards were dynamically allocated and de-allocated to support one or more hypotheses; one blackboard per hypothesis under consideration by M. The controller used several different analytic machine learning techniques to learn management strategies to provide ordered access to the blackboard by the agents.

Given the combinatorics of user behaviors along with the expressive grammar of legal actions/objects in a VMR, M's architecture of modal agents was an effective solution providing two advantages. First, leveraging collaborating "expert agents" was an effective means to obtain multiple perspectives and hypotheses during problem solving; thus providing an extended, expressive range of reasoning ability. Second, each agent was designed leveraging the most expressive knowledge representation(s) (KRs) and algorithms appropriate for its respective modal reasoning. Examples of KR algorithmic treatment included: rule-based, frame-based, case-based, semantic networks, scripting, and numerous analytic learning techniques.

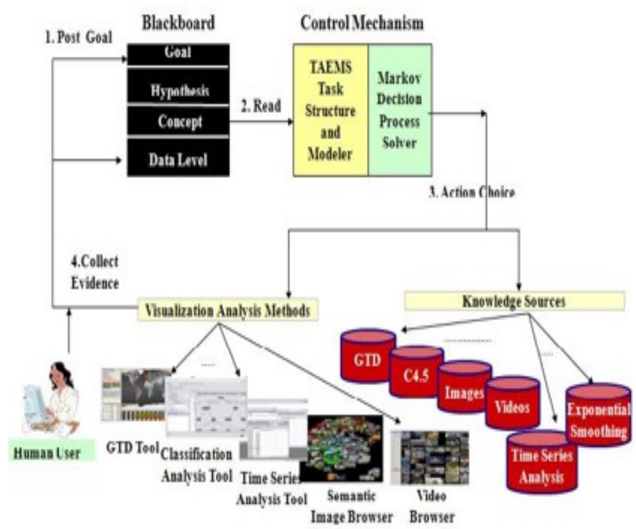

Figure 2: RESIN Control Flow

In prior work such as the Wolfgang system [Riecken, 1989] that reasoned and learned to compose music and the M system, these systems were based on centralized control to manage the diversity of expert agents. In our current work we are interested to enable collaboration of agents via a decentralized technique; thus allowing the agents to learn to collaborate based on context and reward mechanisms.

Evolving From Centralized Blackboards to DEC-POMDPs

In our previous work [Yue, 2009], we designed and developed a visual analytic RESource bounded INformation gathering agent (RESIN) to facilitate access of investigative analysts to (a) automated support for predictive analytics and decision making; (b) the capability for finding non-myopic alternate solution paths; and (c) a tool to investigate outliers. Figure 2 describes the control flow of RESIN.

RESIN harnesses a blackboard to maintain a clear evidential path for supporting and contradicting information while allowing for explicit modeling of concurrent top-down and bottom-up processing. The blackboard contains reasoning results from processing existing information, which includes raw data, various problem-solving states, partial solutions and current goals. Providing clear explanations in support of the decision making process is critical to gain and maintain the analyst's trust in the system.

In addition to the blackboard, RESIN consists of a TÆMS [Decker, 1993] task structure library, a centralized MDP solver and heterogeneous knowledge sources (KSs). TÆMS are hierarchical abstractions of multiagent problem solving processes that describe alternative ways of accomplishing a desired goal; they represent major problems, the decision points and interactions between problems, but they do not describe the intimate details of each primitive action. The MDP is a probabilistic model, which captures the essence of sequential processes and is used to compute policies that identify, track, validate and reject hypothesis. The TÆMS task structure is translated into a MDP by initializing a state set [Wagner, 2006], identifying the possible actions to determine the optimal action choices, and expanding each possible outcome which is characterized by discrete quality, cost and duration values [Liu, 2007], The KSs are independent specialist computational modules that include visual analytics databases and tools of varying complexities that contain the domain knowledge needed to solve a problem. The agent control handles several decisions including selecting databases of high dimensional data for analysis, determining appropriate interactive visualizations for these data, reasoning about the problem solving process; to generate and validate hypothesis.

In SNARE, we will harness our experience in building multiagent architectures with a centralized decision theoretic control mechanism to design a multiagent framework with decentralized control.

A *DEC-POMDP* [Bernstein, 2002] is a tuple: $M = S, A, P, R, \Omega, O$, where

- S is a set of states with a designated initial state distribution $b0$
- _$A = A_1$ X ... X A_n specifies each agent's finite set of actions
- $P(s' | s,a)$ is the probabilistic state transition model
- $R(s,a)$ is the reward model
- $\Omega = \Omega_1$ X ... X Ω_n specifies each agent's finite set of observations
- $O(o|s',a)$ is the stochastic observation model

In the DEC-POMDP model, every agent chooses at each step an action based on its local observation histories (or internal belief state), resulting in a stochastic change in the state of the system, an immediate reward, and an observation for each agent. The true state of the system is unknown and histories of observations represent the knowledge agents have about their environment (including other agents). Reward functions capture both the utility of achieving a task as well as the costs associated intermediate actions. Solutions to DEC-POMDPs are a set of policies where each policy dictates to the corresponding agent the action choice -based on its local observations. A DEC-MDP is a DEC-POMDP with joint full observability. The complexity of DEC-POMDPs and DEC-MDPs for a multiagent system with just 2 agents has been shown to lie in the NEXP-Complete space. The complexity decreases to NP-Complete when joint full-observability as well as transition and observation independence is assumed. If there is single global goal or many decomposable global-goals, the problem complexity reduces to P [Goldman, 2004]. One other way we plan to investigate is to approximate the global reward using a local factored reward function that leverage information about neighbor states; this approach is based on our earlier work in a multiagent weather tracking system [Cheng, 2010].

SNARE Framework for Anytime Situational Awareness

Obtaining anytime situational awareness in a social media space encompasses identifying a compelling range of contextual features, cues, and cultural content. The identification of context for a situation of interest is embedded in a complex fluid space of features and data. This suggests that a range of methods and strategies are in order to obtain multiple perspectives and identify situations of interest. For example, identifying a given situation could require detecting a set of objects/events of interest by applying SNARE agents that collaborate along with human analysts; each agent providing a specialized reasoning ability such as: (1) evidence identification, (2) temporal classification, (3) threat detection, (4) topic modeling, (5) topic propagation, (6) sentiment prediction, and (7) hypothesis explanation generation.

Each SNARE agent is a complex expert in a given domain such as topic modeling or rumor detection or a particular modality of reasoning such as functional or causal reasoning. As each agent reasons and learns, it uses a DEC-POMDP model as its control mechanism to manage its collaboration with other agents.

Cognitive Model of Human Information Analysts

Pirolli and Card [Pirolli, 1999] developed a general model of cognitive task analysis as performed by human analysts.

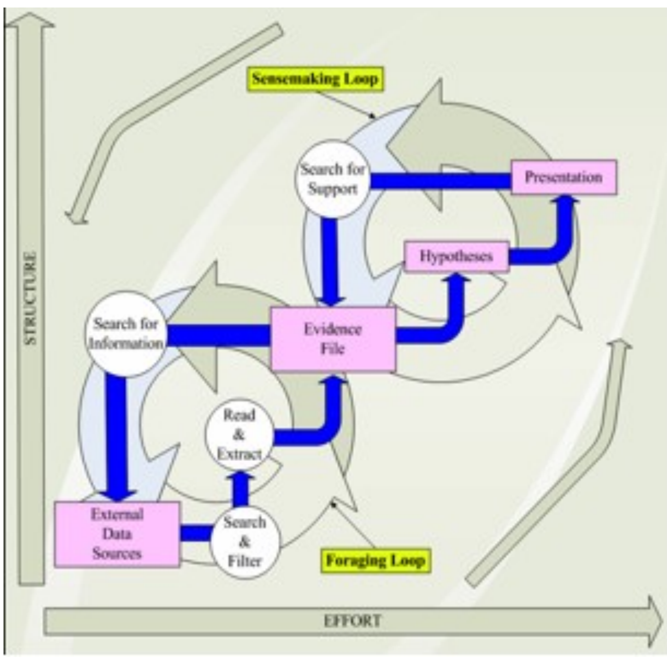

Figure 3: Pirolli and Card Task Analysis

They identified two main, overlapping loops in the analyst's problem solving approach, a foraging loop and a sensemaking loop. Figure 3 depicts this process. In the foraging loop, three processes interplay: exploring (monitoring), enriching (narrowing), and exploiting (more thorough analysis). The sensemaking loop involves hypothesis building, the maintaining of competing hypotheses, and the construction and passing of reasoning artifacts between components of the analysis and hypothesis-building structure between system and users. Assumptions and evidence gathered from the foraging phase, which are frequently highly organized, are used to build hypotheses; attempts at confirmation subsequently reshape evidence gathering and the hypotheses themselves.

The Pirolli and Card cognitive model fits well with SNARE's task process model of collaborating SNARE agents. SNARE is designed to perform monitoring of complex environments similar in fashion to a human information analyst.

Problem solving in the social media space is complex because the problem context and data are constantly changing. Moreover, the data is massive and is typically incomplete, unreliable and/or conflicting. This implies there is a need for deciding which data sources to query, and what types of analysis to use for collecting, assimilating and abstracting the data into evidence. Further, the analysis tasks are usually time critical.

SNARE Architecture

Based on the Pirolli and Card cognitive task model, the SNARE architecture divides its agents into two functional groups; Data specialists (representing the cognitive foraging loop) and Hypothesis specialists (representing the cognitive sensemaking loop). Figure 4 depicts SNARE's architecture. One of our design goals is for SNARE to detect such situations as: user specified topics, deception and misinformation, campaigns intended to persuade and influence, and various threats.

Snare's architecture of specialized Data and Hypothesis agents consists of the following elements:

(1) A set of Data agents, each with expertise in directly accessing and processing social media content (e.g., blogs, Twitter feeds, Facebook, etc…) and performing a specific task (e.g., topic modeling/drift detection, group detection, deception detection, sentiment detection, etc…). Each agent's performance to collaborate, reason and learn is managed by a DEC-POMDP control mechanism and each agent will communicate with other agents via SNARE's Hypothesis Space.

(2) The Hypothesis Space is a collection of data structures/knowledge representations (e.g., "blackboards", "frames", etc…) to enable sharing of information between the all agents (and human analysts) as they collaborate.

(3) A set of Hypothesis agents generate and track hypothesis. They serve to apply multi-strategy reasoning functions to evidence and information in the Hypothesis Space. Unlike the Data agents that access social media resources, Hypothesis agents do not. Like the Data agents, each Hypotheses agent's performance to collaborate, reason and learn is managed by a DEC-POMDP control

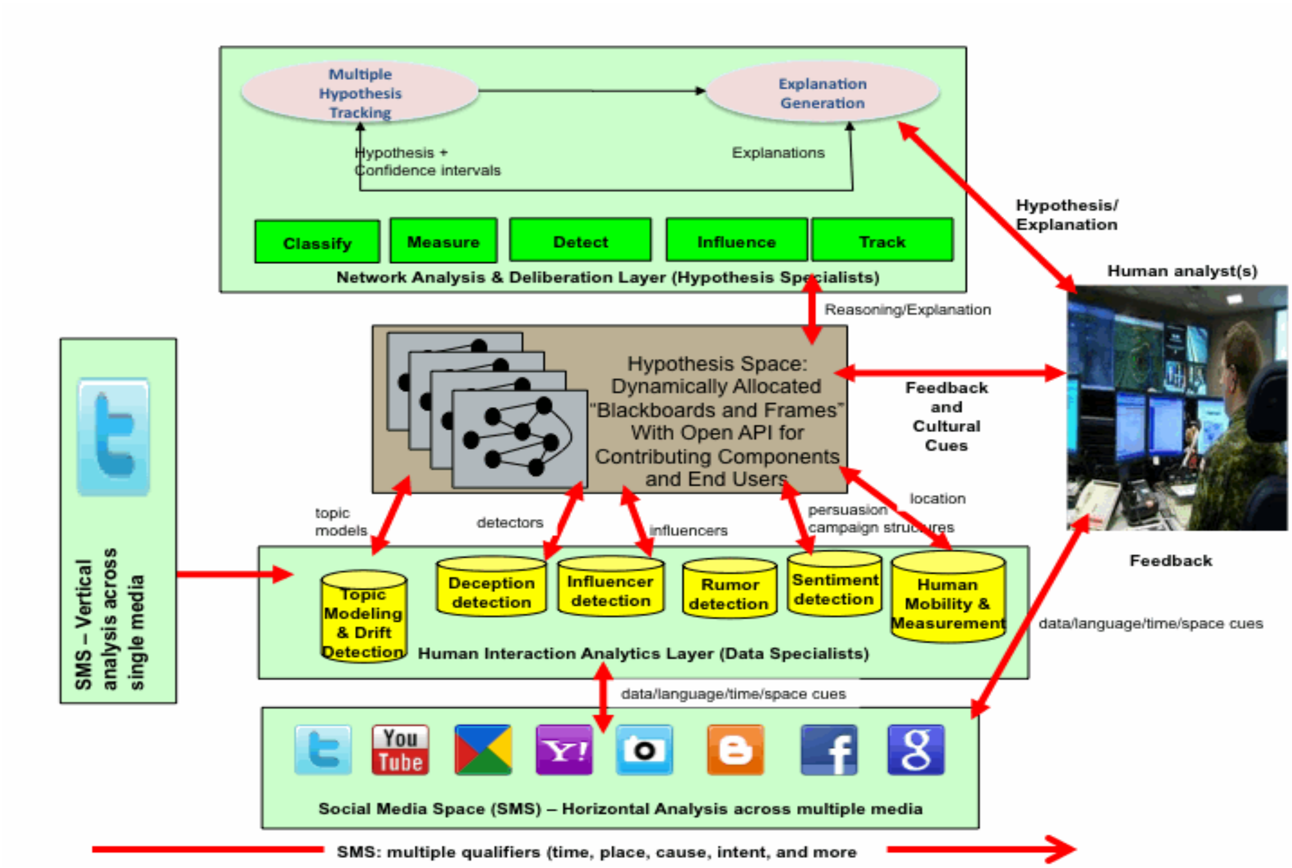

Figure 4: SNARE Architecture of specialized agents

mechanism and each agent will communicate with other agents via SNARE's Hypothesis Space.

(4) An open API that enables read-write access into the shared Hypothesis Space for the collaborating (Data and Hypothesis) agents and applications supporting the human analysts.

(5) Our framework enables human analysts to operate in real-time with the collaborating SNARE agents as social media situations are emerging and being examined. An important goal of this design is to investigate applying a model of mixed-initiative interaction within our framework. Human analysts will be able to collaborate, annotate and influence the real-time processing of SNARE's agents.

We wish to underscore that each SNARE agent can typically be a sophisticated expert agent. For example, several of the Data agents that forage specialize in:

- detecting the spread of ideas and concepts
  - detecting, measuring, and tracking emerging patterns
  - identifying network patterns of interest
- socially conditioned topic modeling
- rumor detection
- deception detection
- persuasion detection
- participant identification
  - human mobility and measurement

As we progress our work, we are focused on DEC-POMDP performance for reasoning and learning; meaning what baseline characteristics are achieved and what improvements can be identified. Also, we look to progress a mixed-initiative model of bi-lateral collaboration between humans and agents.

Looking Ahead

The Role of Minsky's K-lines in Dynamic Agent Organization

We now address the question of how to dynamically form optimal sub-networks of SNARE agents to handle an emerging social media space problem. We design a mechanism that is similar to the formation of knowledge–lines (K-lines) in Minsky's theories of the human mind [Minsky, 1986]. A K-line is defined as a wire-like structure that attaches itself to whichever mental agents and/or groups of agents that are active when you solve a problem. When the K-line is activated later, the agents associated with it are awakened leading a mental state that is similar to the one the agents were in when the problem was solved previously.

In SNARE, each agent will trigger when certain key state features in their DEC-POMDP are assigned particular values called "trigger" values. When a social media space scenario is encountered for the very first time, the scenario would result in certain scenario-specific feature values being set. In SNARE, each agent will trigger when certain key state features in their DEC-POMDP are assigned particular values called "trigger" values. When a social media space scenario is encountered for the very first time, the scenario would result in certain scenario-specific feature values being set. The agents whose trigger values overlap with these scenario-specific feature values will be awakened and will form a problem solving agent sub-network. The formation of the sub-network will then trigger a K-line agent that is looking for this particular combination of agents to be awakened; thus, a new K-line is formed! This K-line and trigger values combination will be stored in memory and will be reactivated the next time the same set of features are triggered.

We suspect this line of investigation will provide opportunities to grow societies of diverse agent communities that are dynamically interacting within and across societies of agents forming heterarchical topologies. This would provide opportunities to examine and measure such concepts as Level Bands and Transframes as described in Minsky's work.

Explanation Generator

Explanation generation in an analytical agent must be autonomous; it should include justifications and explanations of the conclusions and decision-making process and knowledge; and be tunable to the different goals of various human customers. An agent's self-explanation describes the agent's reasoning in reaching a conclusion. For example, in investigative analysis, an automated assistant may provide a description of its decision-making process. A self-explanation in general may have three components [Chandrasekaran, 1991]: (1) justification of conclusions; (2) explanation of the decision-making process; and (3) justification of the decision-making knowledge. In investigative analysis, for example, the evidence for and against it may justify a conclusion about a specific pattern of activity; the decision-making process may be explained in terms of the steps of the process; and the decision-making knowledge may be justified in terms of past cases of investigative analysis. The decision-making process can be explained in many ways, each with its own benefits and drawbacks. We plan to leverage our prior experience in developing explanations of decision-making in the context of intelligence reports [Goel, 2009; Raja, 2007] to the social media space.

Threat-level Manager

We now discuss an outline for a smart, common-sense detection methodology for seeking and validating or rejecting hypotheses that we plan to pursue in parallel to above-mentioned methods. Specifically we discuss hypotheses about threats in the following paragraphs and will extend these ideas to other types of hypotheses about persuasion campaigns, influence operations and deception tracking.

A sound procedure for seeking valid threats is a nontrivial exercise. Johnson [Johnson, 2010] points out how close to detecting the intersection between one pilot seeking to learning to fly a jumbo jet without learning to land, and the approximately twenty students seeking similar training. If this pattern had been detected, 9/11 could have been foiled. Of course, no matter where one draws the line on searches for threats, valid threats may lie just outside a candidate date. But a smart set of policies can minimize missed threats.

First, the key problem for detection of threats is pairs of items that match or are unusual or marked in some way. Single items --- unless totally explicit about their threat --- are seldom a factor. Triples of items are of possible interest, but only after some initial screening. Second, after finding pairs of interest, we need to validate the potential threat, and choose whether to keep the item active. Whenever possible, we plan to use machine learning to improve screening coverage, and look for opportunities for continual improvement for detection and matching of threats. Dependence upon human memory and reason is inadequate and wasteful of resources because humans are likely to forget or miss important items.

Third, we need tools for delineating and then winnowing down a set of candidates to a manageable number of hypotheses to consider seriously. Using information theoretic measures and various means for generating sets for comparison, we can automatically generate and order candidate sets of hypotheses for significant matches. (This process might, of course, involve successively enlarging and winnowing, or other processes.) The processing should have a natural pace of progress, so that the candidacy set for each search step fits in search space with enough processing time for handling a normal processing cycle with all false alarms. Fourth, for candidate threats with sufficient interest[1], we would initiate other steps, ending, potentially in a custom, deep and reasoned search, or ending in the elimination of candidates as false alarms.

Fifth, we need criteria for forgetting items, and heuristics for dropping items from the lists of potential threats. Items should be pruned based on the prominence or significance of the threat, the recency of the threat, the number of connections or intersections between threats, connections based on short chains of inference, etc.

Conclusions

To conclude, we described SNARE, a multiagent framework that has a diversity in reasoning and learning. We discussed the intellectual challenges we have encountered and the progress we have made in the design and development of this framework.

References

Bernstein, D., Givan, R., Immerman, N., and Zilberstein, S., 2002. The complexity of decentralized control of Markov decision processes. *Mathematics of Operations Research*, 27:819–840.

Bertesekas, D. and Tsitsiklis, J., 2006. Neuro-Dynamic Programming. *Athena Scientific*, Belmont, MA.

Chandrasekaran, B. and Swartout, W., 1991 Explanations in Knowledge Systems: The Role of Explicit Representation of Design Knowledge. IEEE Expert 6(3) 47-49, 1991.

Cheng, S., Raja, A., and Lesser, V., 2010. Multiagent Meta-level Control for a Network of Weather Radars. In *Proceedings of 2010 IEEE/ WIC/ ACM International Conference on Intelligent Agent Technology (*IAT-2010) , pp 157-164, Toronto, Canada.

Decker, K. and Lesser, V., 1993. Quantitative modeling of complex environments, *International Journal of Intelligent Systems in Accounting, Finance, and Management*, December 1993, 2(4): pp.215-234.

Goel, A., Morse, E., Raja, A., Scholtz, J., and Stasko, J., 2009. Introspective Self-Explanations for Report Generation in Intelligence Analysis . Proceedings of IJCAI-2009 workshop on Explanation-Aware Computing, Pasadena, California, July 11-12, 2009.

Goldman, C., and Zilberstein, S., 2004. Decentralized control of cooperative systems: Categorization and complexity analysis. *Journal of Artificial Intelligence Research*, 22:143–174.

Johnson, S., *Where Good Ideas Come From: The Natural History of Innovation*, Riverhead Books, New York, 2010

Liu, D., Raja, A., and Vaidyanath, J., 2007. TIBOR: A Resource-bounded Information Foraging Agent for Visual Analytics. In *Proceedings of 2007 IEEE/ WIC/ ACM International Conference on Intelligent Agent Technology (IAT).*

McCarthy, J., Minsky, M., Sloman, A., Gong, L., Lau, T., Morgenstern, L., Mueller, E.T., Riecken, D., Singh, M., and Singh, P., 2002. An Architecture of Diversity for Commonsense Reasoning. *IBM Systems Journal*, 41(3), 530-539.

Minsky, M., 1986. *The Society of Mind*. New York: Simon and Schuster.

Pirolli, P., and Card, S., 1999. Information foraging. *Psychological Review*, 106, 643-675.

Raja, A., and Goel, A., Introspective Self-Explanation in Analytical Agents, Proceedings of AAMAS 2007 Workshop on Metareasoning in Agent-based Systems, pp 76-91, Hawaii, May 2007.

Riecken, D., 1994. M: An Architecture of Integrated Agents. Special Issue of the *Communications of the ACM - Intelligent Agents*, D. Riecken (ed.), July 1994, Vol. 37, No. 7. ACM: New York, New York.

Riecken, D., 1989. Goal Formulation with Emotional Constraints: Musical Composition by Emotional Computation. In *AAAI Proceedings First Annual Conference on Innovative Applications of Artificial Intelligence*, Stanford University, AAAI/MIT Press: Cambridge, Massachusetts.

Wagner, T., Raja, A., Lesser, V., 2006. Modeling Uncertainty and its Implications to Sophisticated Control in TAEMS Agents" Autonomous Agents and Multi-Agent Systems, Vol 13, Number 3, November 2006, PP 235-292.

Yue, J., Raja, A., Liu, D., Wang, X., and Ribarsky, W., 2009. A Blackboard-based Approach towards Predictive Analytics. *Proceedings of AAAI Spring Symposium on Technosocial Predictive Analytics*, pp 154-161, Stanford University, CA, March 23-25, 2009.

[1] There are likely good heuristics at this level. For example, if any items that make it to stage four make a short list and if these items ever show up in others parts of the search, the significance would be increased and noted.